Postel

Network Working Group
Request for Comments #323
NIC #9630

V. Cerf
UCLA-NMC
23 March 72

Recd 7APR72

Formation of Network Measurement Group (NMG)

On March 17, 1972, at MIT project MAC, the following group met to discuss plans to perform measurement experiments on the ARPANET:

A. Bhushan - MIT/DMCG
V. Cerf - UCLA/NMC, Chairman, NMG
S. Crocker - ARPA/IPT
J. Forgie - LL/TX-2
R. Metcalfe - MIT/HARV/DMCG
M. Padlipsky - MIT/MULTICS
J. Postel - UCLA/NMC
J. Winett - LL/67

The purpose of the meeting was to discuss existing and planned measurements of network and HOST behavior.

1. Measurement Link #'s

It was agreed (after a ridiculously long discussion) to allocate links 159-191 for network measurement only (see RFC #317). It was further agreed that these links would be allocated in the following way:

159-174	HOST DISCARD; co-operating HOSTS receiving messages on these links will throw them away *without* generating an error message.
175-190	To be allocated as needed by V. Cerf - UCLA/NMC.
191	To be used by IMPs to send measurement traffic obtained from IMP statistics packages.

It should be apparent that HOSTs wishing to co-operate in the support of a HOST discard service should modify their NCP's to throw away all messages

received on links 159-174 without sending an error back to the source HOST (no connection will be open on these links).

2. Process Discard

Although it was not mentioned at the meeting, C. Kline at UCLA has suggested a PROCESS DISCARD also with some well known socket number. The purpose of this discard routine would be to help us study Process-Process behavior of the network.

It would be convenient if all co-operating HOSTs could write a Process Discard program which would simply wait for ICP on some standard socket number. Until a complete survey is made of well-known socket numbers at each HOST, no socket number will be proposed (see RFC #322).

3. NCP Statistics

At the meeting it was apparent that several sites have already instrumented their NPC's out of curiosity. In particular, Joel Winett, Lincoln Labs (360/67), has instrumented all connections originated by local TELNET users. He gathers statistics per connection such as:

a) Network connect time
b) NCP CPU time
c) Number of reads or writes on connection
d) Time stamps on:
    first RFC, last RFC, first close, last close.
e) Number of messages and bits transmitted
f) Log of errors sent or received

MULTICS gathers summary statistics on the number of regular (type 0) messages sent and received, and the number of irregular messages (not type 0) sent or received.

The NWG agreed to implement a minimal NCP instrumentation procedure which would gather by HOST for some standard 24 hour period (e.g. local midnight to local midnight) the following:

a. Total bits sent to HOST
b. Total bits received from HOST
c. Total messages sent to HOST
d. Total messages received from HOST and optionally
e. Average Round Trip delay on send connections to HOST

The information above should be collected only for standard open connections (i.e. those using standard NCP protocol) and not measurement links or experimental NCP links, and in particular, not traffic on link 0).

Another optional measurement would be to gather the distribution of message types over link 0 over all HOSTS (i.e. not broken down by HOST). This will reveal the relative utilization of control messages (ALLOC should be very prevalent).

The data collected for the last 24 hour sample period should be available from a process whose well-known (to be specified) socket number will support ICP and will produce a message in the following format:

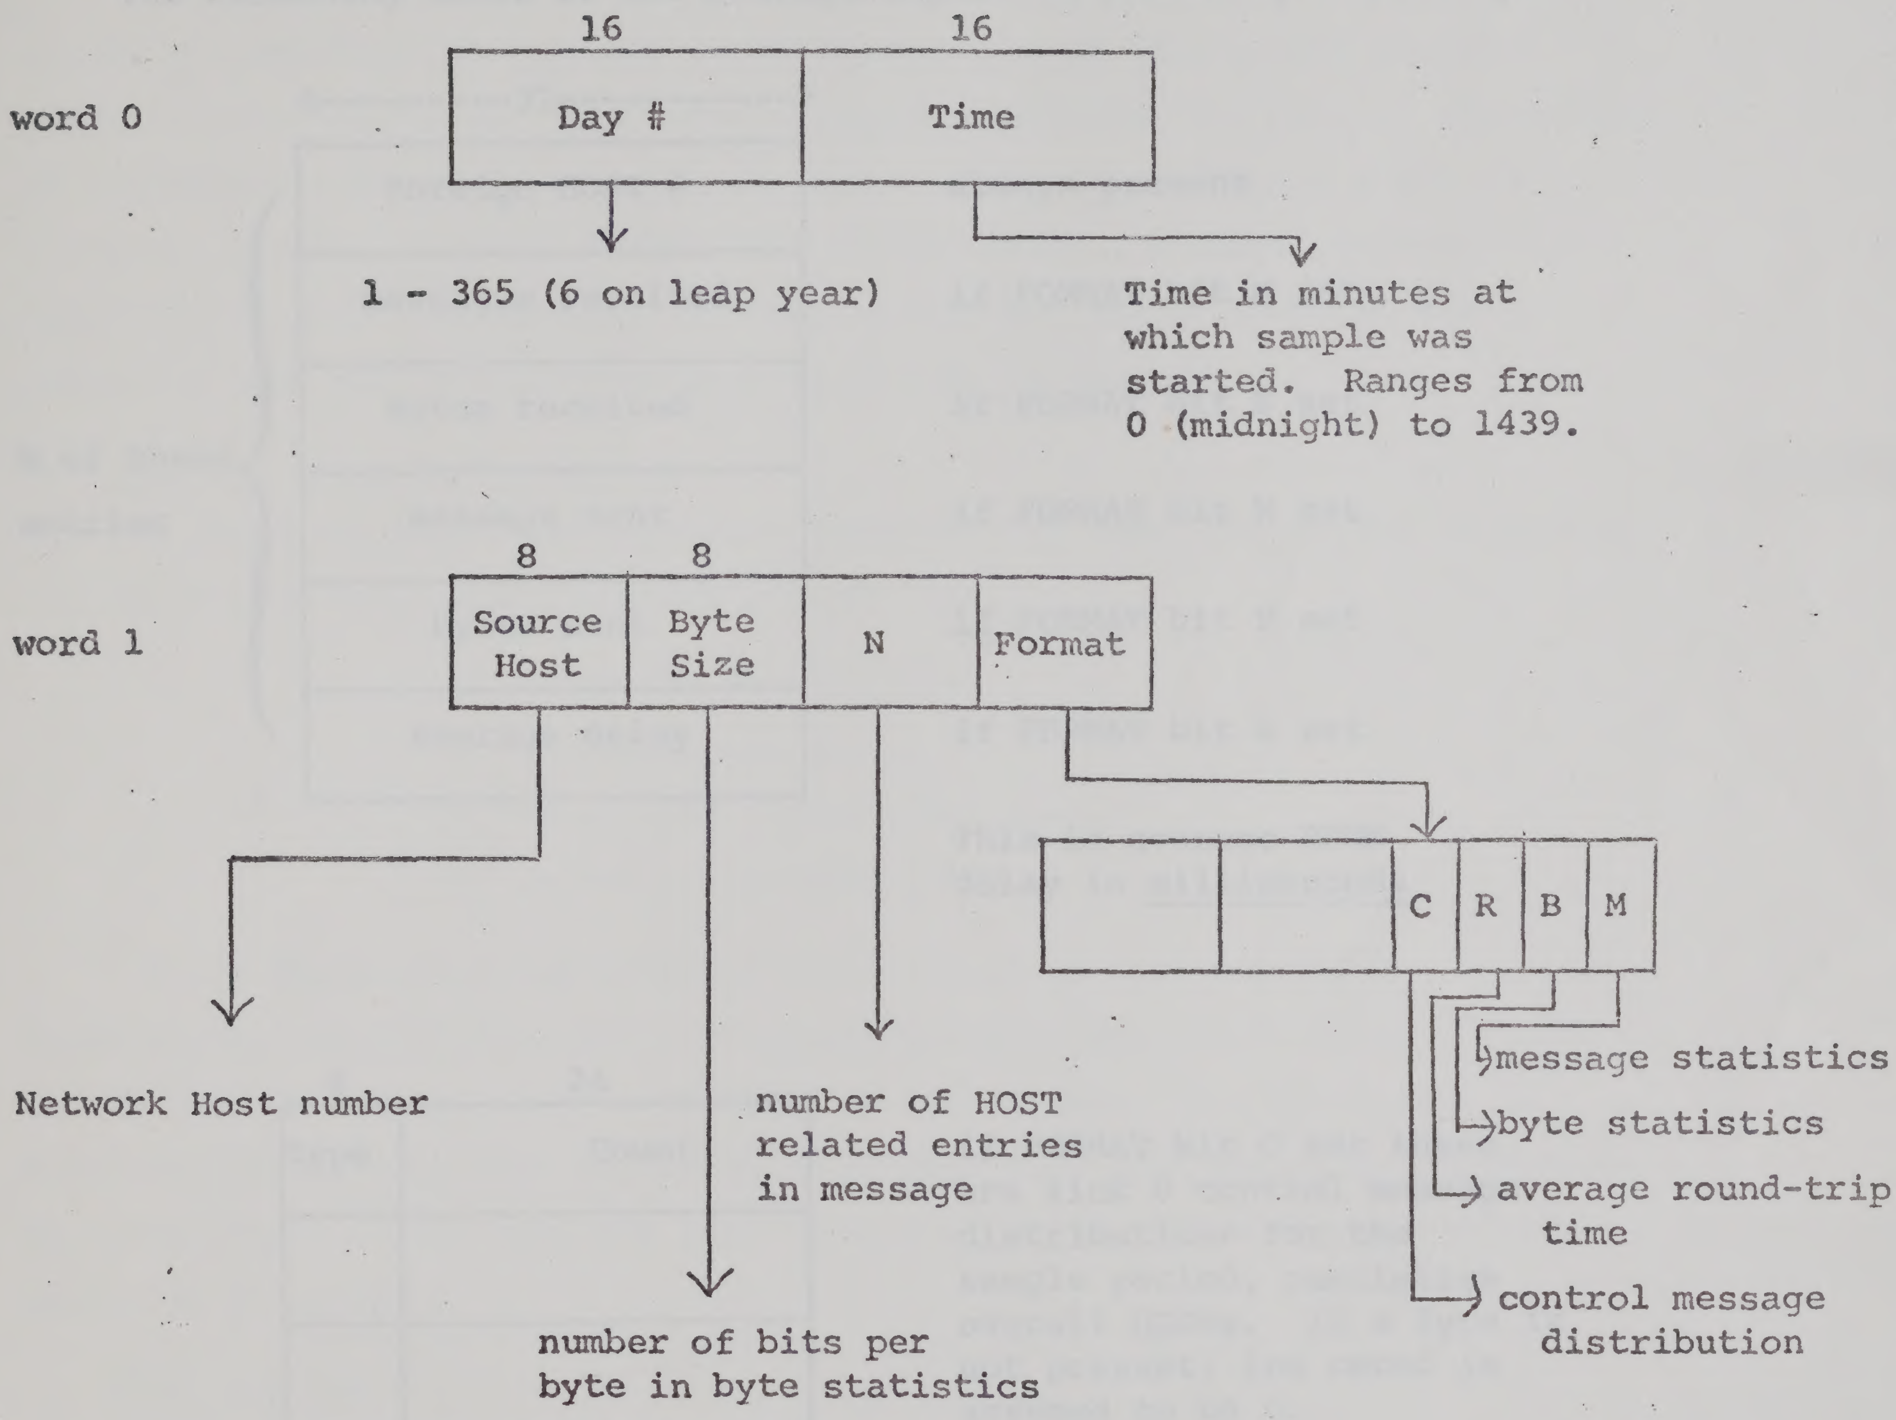

Time in minutes at which sample was started. Ranges from 0 (midnight) to 1439.

The remaining words of the message depend on Format byte settings:

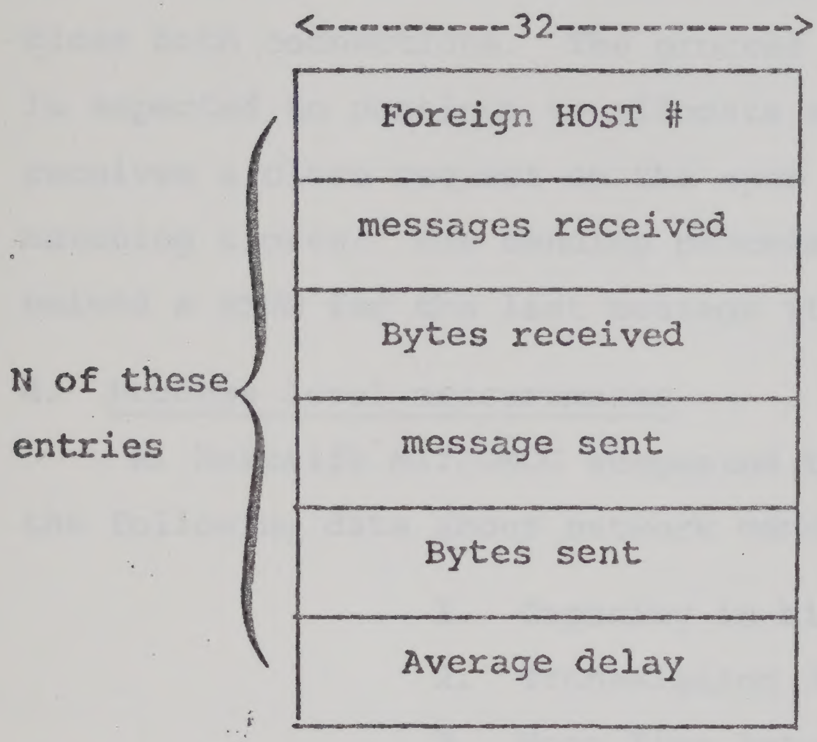

always present
if FORMAT bit M set
if FORMAT bit B set
if FORMAT bit M set
if FORMAT bit B set
if FORMAT bit R set

This is average RFNM delay in milliseconds

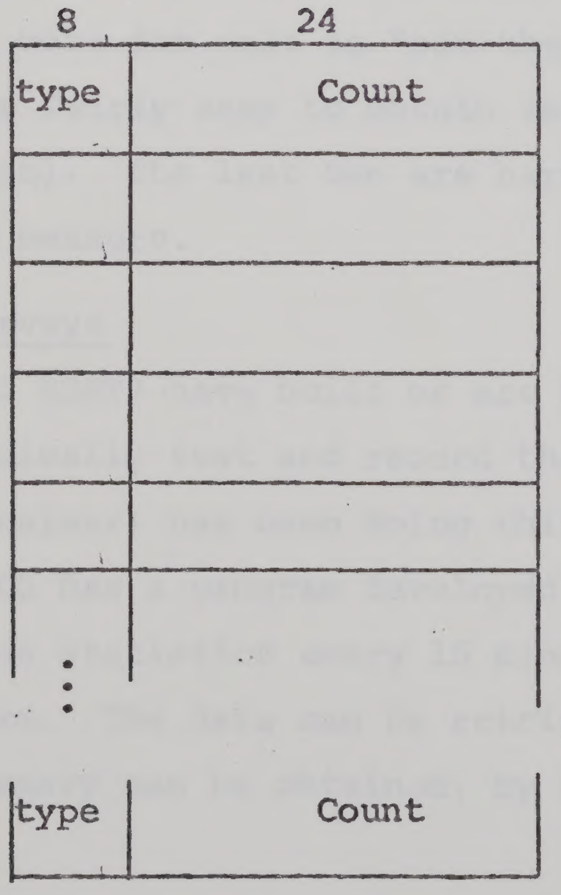

if FORMAT bit C set these are link 0 control message distributions for the sample period, cumulative overall HOSTs. If a Type is not present, its count is assumed to be 0.

The process sending these statistics will continue to send data until it has transmitted the entire statistics sample at which time it will close both connections. The process which requested the initial connection is expected to continue to allocate space as it is available until it receives a close request on the open connections. It then responds with matching closes. The sending process should not close until it has received a RFNM for the last message it wishes to send.

4. Process level measurements

R. Metcalfe MIT/DMCG suggested that the NMG consider trying to gather the following data about network connections:

1. Capacity in bits/sec
2. Transmission delay
3. Mean Time Between Failures
4. Percent availability

These statistics characterize connections as communication commodities and would be the kind of information one would want if Network connections were for sale as "off-the-shelf" items. The first two measures are fairly easy to obtain (although they may vary from connection to connection). The last two are harder to get at and will require some planning to measure.

5. HOST surveys

Several HOSTS have built or are building automatic survey programs which periodically test and record the status of various HOSTs. BB&N (Ellen Westheimer) has been doing this manually on a daily basis.

MIT/DMCG has a program developed by R. Metcalfe and M. Seriff which gathers these statistics every 15 minutes and stores the data away in massaged form. The data can be retrieved through the NETWRK program at DMCG. A summary can be obtained, by HOST, declaring the % time VP overall

samples and the message response to perform ICP in seconds. This program also keeps the state of the HOSTs according to the following measures:

code	meaning
0	HOST not surveyed
1	HOST Dead (according to IMP)
2	NCP not responding to RESET request (15 second time-out)
3	NCP rejecting (ICP got close response).
4	Logger not responding (20 second time-out after ICP request).
5	Logger available (i.e. ICP successful followed by Close request by DMCG).

Details and sample data are available in an RFC produced by M. Seriff (RFC #308, NIC #9259). At UCLA, M. Kampe is implementing a similar program.

J. Postel and V. Cerf plotted Ellen Westheimer's data for HOSTS OPEN (disregarding HOST advertising of service hours) and found the resulting plot rather interesting. The result is reporduced in the figure below. On a moving average, the number of HOSTS OPEN seems to be increasing, which is a good sign.

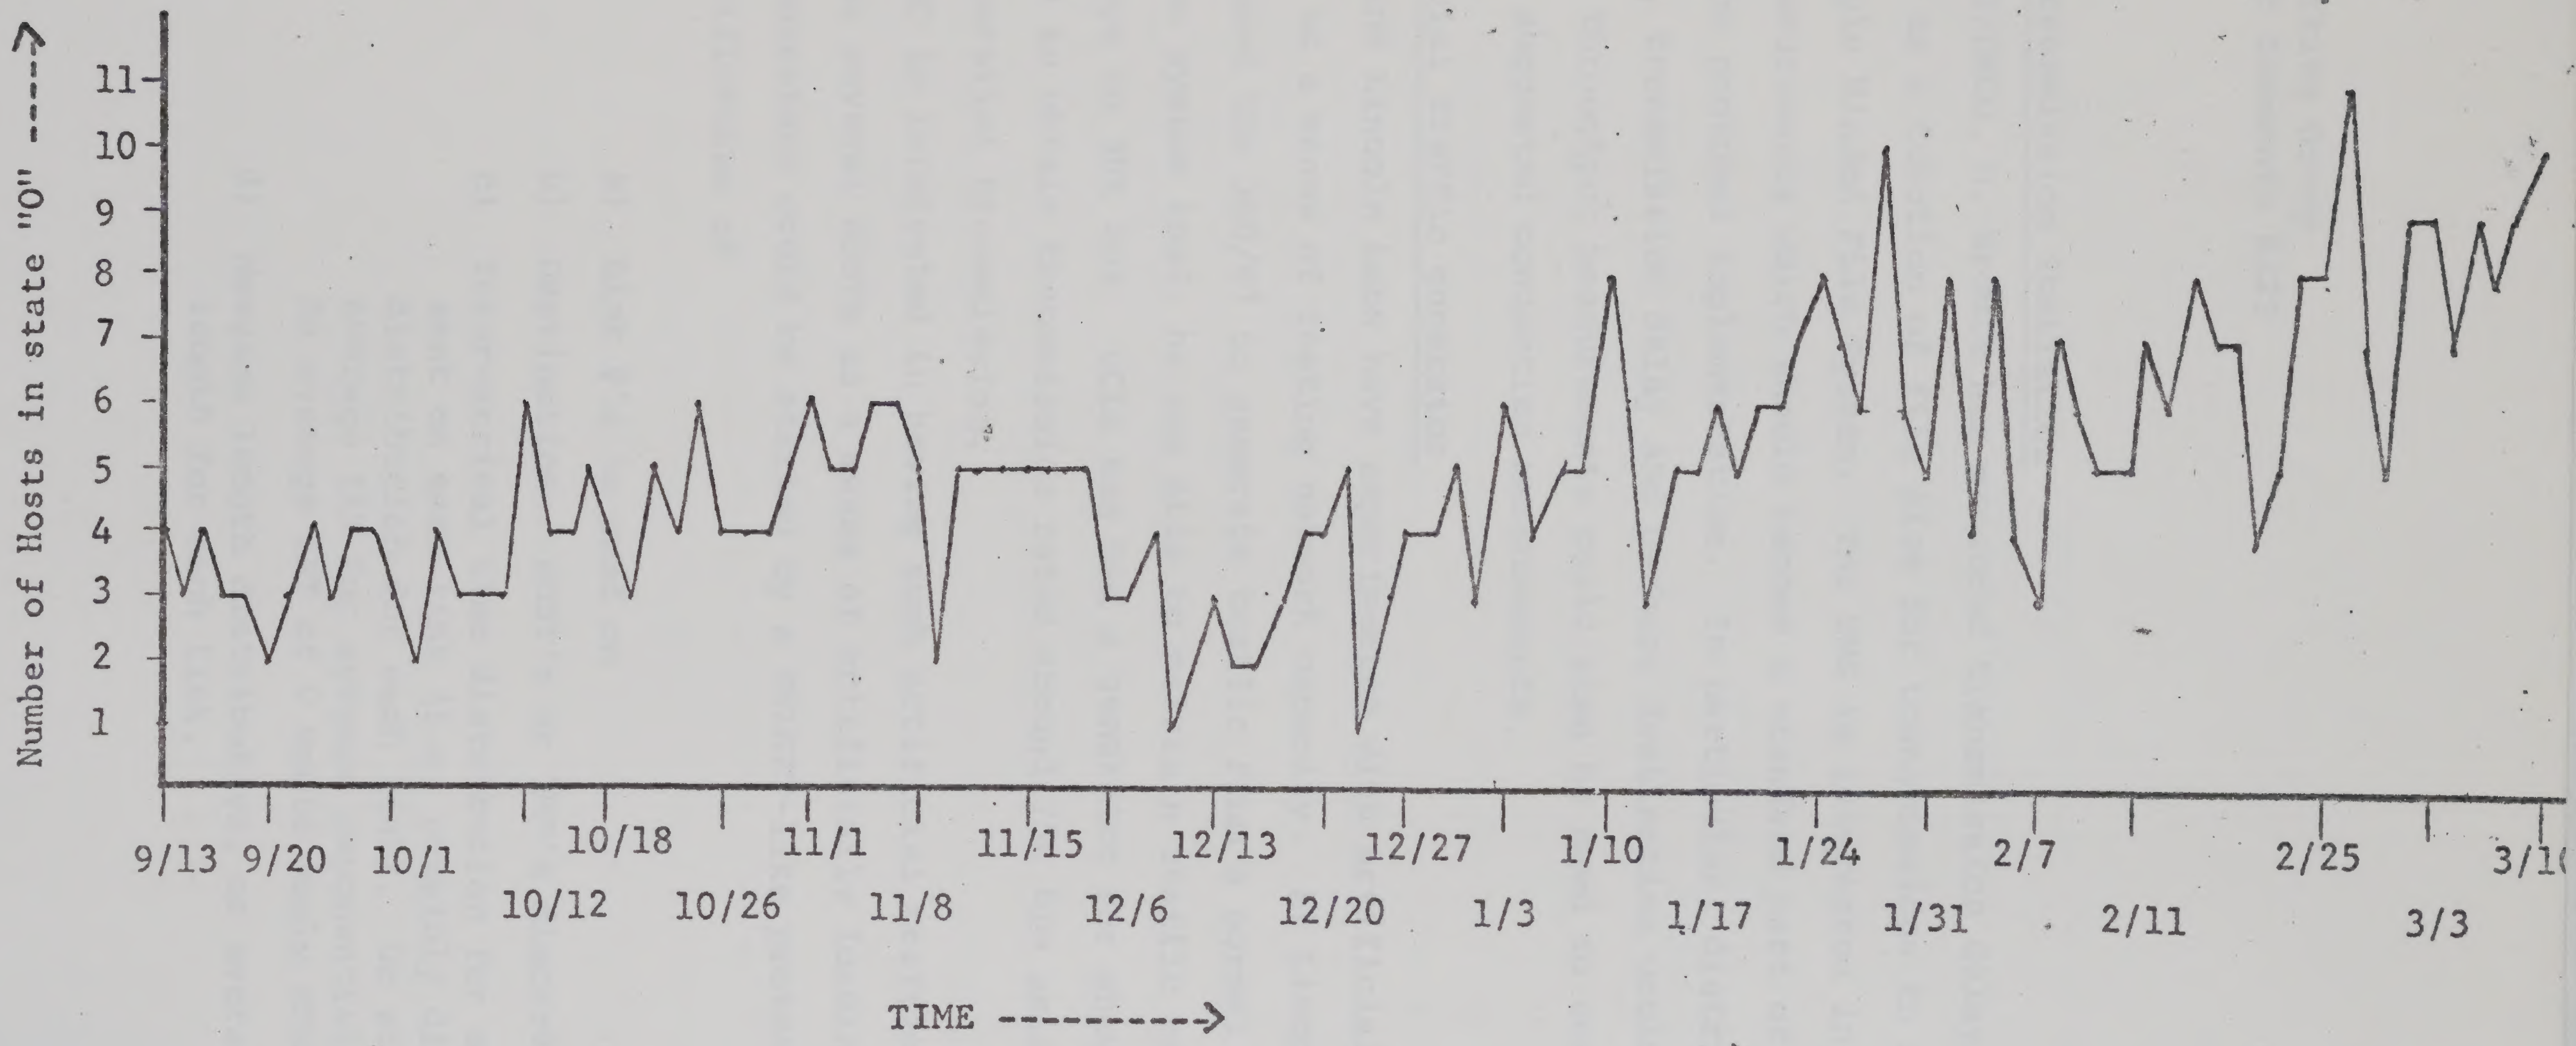
Number of Hosts in state "0" ---->
1 2 3 4 5 6 7 8 9 10 11
9/13 9/20 10/1 10/12 10/18 10/26 11/1 11/8 11/15 12/6 12/13 12/20 12/27 1/3 1/10 1/17 1/24 1/31 2/7 2/11 2/25 3/3 3/1
TIME --------->

5. File Transmission Statistics

At MIT/DMCG, H. Brodie has measured transmission delay and total throughput as a function of file size for transmissions to and from UCSB's Simple Minded File System. The NMG is interested in specifying certain measurements which should become a standard part of any File transmission protocol implementation. In particular, distributions of file sizes, transmission delay and perhaps destination would be of interest. Throughput measurements could also be used to correlate with Metcalfe's suggested connection measurements.

6. Artificial traffic generator

UCLA and Lincoln Labs have experimented with artificial traffic generators as a means of testing network capacity. At Lincoln Labs, J. Forgie used the 360/67 to generate traffic from a normal user process. Depending on system load, he was able to maintain traffic rates ranging from 4800 bps to 38K bps. UCLA has had a generator for about a year and has managed to obtain transmission rates around 75K bps using multiple links for parallel transmission.

The NMC is interested in having such artificial traffic generators available at several HOSTs as a means of artificially loading the network. Ideally, generators could be started by a TELNET-like protocol and would permit specification of

a) Link #'s to send on

b) Destination: HOST's or IMP's discard

c) Inter-arrival time distribution for messages sent on each link (i.e. possibly different distribution for each link). Or at least average IAT for assumed exponential distribution. An average IAT of 0 would imply RFNM driven traffic

d) Message length distribution, or average, or fixed length for each link.

It would also be helpful to accumulate average round-trip times and total bits sent for the duration of the experiment.

At UCLA, the traffic generator permits the following specifications:

a) message header (includes link #)

b) message length (for each link) - distribution (can be constant for each link)

c) message inter-arrival time - distribution for each link

d) Duration of generation in seconds

We can also send imperative commands to the program to stop message generation prematurely. Throughput and average response times (Round Trip delays) are automatically accumulated for each link and are published at the end of the experiment.

A more sophisticated version will also permit specification of ICP socket number for the Process Discard experiments. The idea is to have a number of artificial traffic generators available at different sites and to be able to start these up remotely from UCLA/NMC during the course of a measurement experiment. More details of the desired generator will be published in another RFC.

7. Measurements at other sites

People at sites not mentioned may have done some measurement work and the NMG encourages these people to publish their results. If anyone is is interested in co-operating with the NMG in making NCP measurements (or what-have-you), please get in touch with

Vint Cerf
UCLA-NMC Computer Science Department
3804 Boelter Hall
Los Angeles, California 90024

(213) 825-4864
(213) 825-2368